

Intel's Itanium:  
Who Benefits from  
Early Adoption?

An Executive White Paper

November 2000

# Intel's Itanium: Who Benefits from Early Adoption?

## Preface

Internet commerce and large database applications are dealing with ever-increasing quantities of data, and demands placed on both server and workstation resources are increasing correspondingly. One demand is for more memory than the 4 GB provided by today's 32-bit computer architectures. But the huge memory addressability provided by 64-bit architectures is only one of the requirements imposed by today's power-hungry applications. These applications also need a processor architecture that provides high performance across a broad range of application types together with a long-term roadmap for future enhancements.

Intel's new Itanium architecture is not the first architecture to offer 64-bit addressability. Current 64-bit RISC (reduced instruction set computer) architectures — most notably Hewlett-Packard's PA-RISC, Sun's UltraSPARC, Compaq's Alpha, and IBM's PowerPC — support 64-bit versions of key applications. But "64-bitness" is just one aspect of a high-end architecture. The processor itself must be designed to provide application throughput through design features that promote efficient use of the processor resources as part of a complete system. Equally important, the processor must be part of a rich infrastructure of hardware components, independent software vendor (ISV) tools and applications, and even firmware.

From Aberdeen's perspective, answering the question of when Information Systems (IS) executives should deploy a given application on the Itanium architecture depends on the careful matching of the features and capabilities of both the Itanium architecture as a whole, as well as those of individual Itanium processors, to IS buyer and application characteristics. The question may be if, not when, IS executives should deploy Intel-based enterprise applications on the Itanium architecture. But is the first Itanium processor the right choice for early adopters? The Itanium processor is expected to begin shipping in end-user pilot systems in late 2000. Intel intends to follow the initial Itanium processor with additional members of the Itanium processor family — with McKinley available for pilot systems in late 2001, and Madison and Deerfield available for pilot systems in 2002. Should IS buyers base high-end enterprise Intel architecture deployments on one of these later processors rather than the initial Itanium processor?

## Executive Summary

The Itanium processor will be Intel's first 64-bit processor, but it represents a new architectural direction rather than a simple extension to Intel's continuing product line of 32-bit processors. This processor is based on a new, explicitly parallel instruction set architecture — an architecture designed to deliver sustainable increases in application throughput by changing the way in which the processor hardware interacts with code.

In this *Executive White Paper*, Aberdeen provides an overview of Intel's Itanium processor program and a detailed discussion of the Itanium architecture.

This paper includes sections that:

- Describe the underlying philosophy and potential of Intel's Itanium architecture;
- Examine the first processor based on the Itanium architecture, as well as its surrounding hardware infrastructure and its competition;
- Analyze the fit between Itanium and different application types with respect to software-driven platform requirements and application deployment characteristics;
- Present specific profiles of IS buyers that can achieve near-term benefits by deploying Itanium-based systems in production environments; and
- Recommend methodologies that IS executives should consider when migrating to the Itanium architecture.

Based on our findings, Aberdeen draws the following conclusions about Itanium, both the architecture and the processor:

- Applications types such as large databases, business intelligence, online transaction processing, e-Commerce solution stacks, security, and technical/scientific stand to benefit from Itanium architecture features such as large memory support, floating-point performance, large register set, and explicitly parallel design.
- The Itanium architecture incorporates hardware-based features for enhanced reliability that are likely to be leveraged by platform vendors to help deliver more robust systems.
- ISVs will likely accept the Itanium architecture as a de facto standard and will migrate their 32-bit software to Intel's Itanium architecture as their applications grow and mature to the point of benefiting from Itanium architectural features such as a high level of instruction-level parallelism (ILP) and very large memory (VLM).
- Over time, systems built around Itanium processors will become major, perhaps leading, volume players in enterprise application deployments — if Intel continues on its current publicly stated schedule — within five to seven years of its debut. Consequently, Aberdeen recommends that IS executives identify an Itanium architecture adoption strategy.
- Aberdeen suggests that a wide range of IS executives consider pilot deployments of Itanium processor-based systems to determine how their key enterprise applications will perform on the new platform — and how to tune these platforms to best fit the applications.
- Early enterprise adopters with a defined set of business needs that align with initial Itanium processor features and solution stack maturity

should consider near-term production deployment of Itanium-processor-based systems.

### **Itanium: The Processors and the Architecture**

Itanium is now the name for the new Intel architecture formerly known as IA-64. Intel extends a branding strategy — begun with Pentium — to encompass not only the first processor based on the new architecture, but also the family of follow-on processors and even the architecture itself. By extending the Itanium brand's umbrella over the architecture as well as the specific processor implementations, Intel is clearly positioning the new architecture as being about much more than just 64 bits. Such positioning is not new, but past "IA-64" terminology inevitably focused attention on 64-bit memory while relegating other new architectural features enhancing scalability, platform longevity, and high availability to secondary roles.

The Itanium architecture's ability to address a flat 64-bit memory address space in the millions of gigabytes has been the focus of attention. Beyond VLM support, however, other traits — including a new Explicitly Parallel Instruction Computing (EPIC) design philosophy that will handle parallel processing in a new way — differentiate Itanium from other architectures. EPIC, developed jointly by Intel and Hewlett-Packard, changes how the microprocessor interacts with software by allowing compilers to specify precisely the manner in which instructions are executed on a processor. This design philosophy is intended to address issues that other architectures — many designed 10 years to 20 years ago — have not yet fully overcome, including achieving maximum ILP and ensuring future growth by simplifying key aspects of the hardware design. Itanium processors also contain massive chip execution resources including registers, functional units, and caches, all of which have the potential to translate into processing power beyond what is available today on other processors. In addition, new reliability features have been added, such as extensive ECC (error correction code) coverage, enhanced machine check architecture, and error logs for FRU (field replaceable unit) identification.

Beyond explicitly parallel code generation, the Itanium architecture incorporates additional features such as predication and speculation intended to address performance-limiting factors — such as branch mispredicts and latency to memory — in current processor designs. Although it has the ability to run IA-32 binaries in a compatibility mode, the Itanium architecture is a completely new processor architecture with a new instruction set, as opposed to an evolutionary enhancement to IA-32/x86.

The appendix of this paper provides in-depth details of the Itanium architecture and the EPIC design philosophy, including the four facets of the Itanium architecture of the greatest interest to IS buyers considering adoption of the architecture:

1. 64-bit pointers and other 64-bit data types;
2. Scalable, explicitly parallel architecture;


3. IA-32 instruction set compatibility mode; and
4. Innovative hardware reliability features.

### Solution Stack Components for Itanium-Based Systems

IS executives anticipating the availability of new processors and new processor architectures often focus on the processors themselves to the exclusion of the surrounding hardware and software infrastructure. Yet, in Aberdeen's experience, this infrastructure often plays the critical role in the usability and value proposition provided by a processor, especially when it is the first implementation of a new architecture.

During the development of the Itanium architecture and processors, Intel has worked closely with key industry suppliers to enable the rollout of solution stack components in parallel with initial and follow-on Itanium processor delivery. This stack includes Itanium-based server hardware, firmware, supporting chip sets, application development tools such as compilers and code migration tools, 64-bit operating systems, and applications that have been rewritten or recompiled for Itanium. Multiple levels of hardware and software components make up the complete Itanium architecture ecosystem (Figure 1), including:

**Figure 1: Itanium Processor Solution Stack Components**



Source: Aberdeen Group, October 2000

- Enabling hardware such as systems, chip sets, firmware, and input/output (I/O);
- Enabling software, including operating systems, compilers, and other software infrastructure components; and
- Application software such as databases and business applications.

### *Enabling Hardware*

Enabling hardware components complement processor technology in the creation of core system platforms and include not only the chip sets, motherboards, and hardware platforms themselves, but also firmware and other base-level abstraction and communication layers.

- All major high-end system vendors — except Sun Microsystems — are committing to using Itanium processors in part of their product lines. Intel architecture vendors like Dell, Compaq, and Unisys, as well as traditional RISC vendors like Hewlett-Packard and IBM, have all stated intentions to deliver Itanium-based systems. Systems planned by these vendors range from quad-processor “commodity” systems to 32-processor systems focused on the largest scale enterprise applications.
- Intel's 460GX chip set supports up to four processors and 64 GB of SDRAM (synchronous dynamic random access memory) to enable quad-processor Itanium-based systems. OEMs (original equipment manufacturers) such as Unisys, NEC, and Hitachi have announced plans to develop higher processor-count Itanium-based systems based on their own chip set designs.
- READY64 firmware from Phoenix Technologies implements the Developer's Interface Guide for Itanium-based systems (DIG64) for both the Intel 460GX chip set and OEM designs.
- The IA-64 Fast Track Driver Development Program is focused on helping independent hardware vendors (IHVs) such as Adaptec, Q-Logic, 3-D Labs, and Matrox enable their I/O devices for Itanium-based systems.

### *Enabling Software*

Enabling software is a broad category that encompasses both tools, such as compilers, necessary for the creation of applications and the multiple layers of software that enable base application functionality.

- Four operating system vendors (OSVs) — Microsoft (Windows 2000), Hewlett-Packard (HP/UX), IBM/SCO (AIX 5L), and Novell (Modesto) — have stated their support for Itanium. In addition, the IA-64 Linux project is well underway with four distributions of development code re-

lease. All are booting versions of their operating systems on Itanium-based software development platforms.

- Compilers from Microsoft, EPC, and IBM are already in use, porting applications from current platforms to Itanium in C<sup>++</sup>, Fortran, Java, and other languages. SGI's Pro64 compilers are open-source performance-oriented tools for C<sup>++</sup> and Fortran Linux development.
- Other components in the Itanium software development environment include code porting tools, such as MigraTEC's MigrationSUITE, and the Cole utility from Microsoft included in the Windows 2000 Software Development Kit (SDK).
- Systems that provide application, Web, and infrastructure services will include cache/Web servers such as the Microsoft Internet Information Server (IIS), Apache, and Inktomi Traffic Server; security products such as RSA CryptoC and Check Point Firewall; and directory servers such as Microsoft Active Directory and Novell Network Directory Services (NDS).
- Middleware and Enterprise Application Integration (EAI) software such as BEA Weblogic and IBM Websphere will serve to integrate software stack components including back-end enterprise applications and e-Business stacks.

### *Application Software*

Application software delivers specific business value to users either through direct interaction or software intermediaries.

- Databases committed for the large memory addressability of the Itanium architecture include Microsoft SQL Server, Oracle 8i, and IBM DB2.
- Business processing applications planned for Itanium include traditional Enterprise Resource Planning (ERP) applications like SAP R/3 in addition to newer e-Business solutions like i2's RHYTHM for Supply Chain Management (SCM).
- Business intelligence applications planned for the Itanium architecture can perform a range of functions, including information management, retrieval, and analysis. For example, IBM's Content Manager provides storage, management, and distribution of digital content while SAS Institute products target complex analysis of many types of data.
- Technical computing applications include financial simulations, digital content creation (DCC), mechanical and electrical design automation (MDA/EDA), and pure scientific applications such as those used for computational chemistry. The large set of applications planned for Itanium in this area include products from vendors such as Riskmetrics

Group (risk management), Alias|Wavefront (DCC), and Fluent (flow and heat transfer modeling).

*Helping to Drive ISV Development of Itanium-Based Solutions: The Intel 64 Fund*  
Intel recognizes the need for complete application solutions for Itanium-based platforms. To foster application development by ISVs, especially in Internet and emerging enterprise spaces, Intel has co-invested with system suppliers and select large Fortune 500 IS executives in an Intel 64 Fund that is being used to help ISVs develop applications for both initial and follow-on Itanium processors. Together, suppliers and IS executives are targeting a portfolio of specific ISVs, bridging the gap between technology development and IS executives who will deploy these application platforms. The fund, valued at \$250 million, complements other initiatives currently underway, including software-enabling programs, application solution centers (ASCs), Intel Itanium-based SDKs, and Enterprise Technology Centers.

*Intelligent Hardware Reliability: Migrating System Robustness to Hardware*  
Enterprises require exceptionally high robustness for business-critical production environments. Today, most large enterprises run business-critical applications on a combination of mainframe and RISC-based Unix server environments, but Intel Architecture servers are starting to grow in importance in enterprise environments. RISC-based server suppliers have been gradually introducing hardware reliability features to their systems and, over time, Intel has enhanced its IA-32 architecture with hardware-based reliability features to reduce system downtime. IS executives have also used other tactics to achieve reliability, including the use of enterprise fault-tolerant servers as well as various forms of high-availability clustering.

The introduction of the first Itanium processor continues this drive toward higher server reliability. To achieve what Aberdeen calls "intelligent hardware reliability," the Itanium architecture employs hardware features to handle error detection, correction, and "containment." These features include enhanced ECC — a more extensive machine-check architecture — improved processor-firmware-OS (operating system) integration, and a host of other improvements. (For more details on these features, see "Itanium Architecture Intelligent Hardware Reliability Features," the last section of the appendix.)

### **When Will Itanium-Based Systems Make Their Debut?**

To date, Intel has delivered more than 6,000 prototype Itanium-based systems to OSVs and ISVs for use as development platforms. The goal is to enable platform and OS qualification simultaneously with targeted end-user pilots during Q4 2000. Beta versions of the OS for the Itanium processor are also available.

- In November 2000, Microsoft announced the beta 1 release of "Whistler," a next-generation version of Microsoft Windows planned for re-



lease on Itanium-based systems as Microsoft Windows 2000 64-bit Edition (Win64).

- Development releases of IA-64 Linux are available as distributions from Red Hat, SuSE, TurboLinux, and Caldera.
- IBM and SCO have shipped a beta version of AIX 5L (formerly Monterey/64) to more than 50 ISVs with a broader beta program scheduled to start in the near future.
- Hewlett-Packard is working with ISVs on the development of database and technical computing stacks and preparing for end-user pilots.
- Novell has the Novell Internet Caching system running on Itanium-based platforms under Monterey, a focused 64-bit operating system for Internet caching and directory services.

Intel expects to continue expanding the large pool of ISVs that have already compiled their applications to run on Itanium-based platform. Working with key vendors, Intel has provided Itanium transition tools to a broad ISV community to assist developers with application ports. Major ISVs tell Aberdeen that they are finding the available tools — such as compilers — to be reasonably robust and functional and, therefore, not an impediment to moving applications to the Itanium-based platforms.

End-user availability of ISV applications will be driven, in large part, by the impact of Itanium processors on a given application. As discussed in the subsequent section, "How Itanium Will Perform on Key Application Types," some application types will benefit more than others from being rewritten and recompiled for the Itanium processor. Current 32-bit commercial applications with modest memory and performance requirements deployed independently of large database applications may see little near-term benefit from moving to the Itanium processor. Such applications may, therefore, remain 32-bit for the foreseeable future. Aberdeen notes that many of the applications deployed on 64-bit RISC processors today have never been rewritten for 64-bit. Rather, these 32-bit applications coexist on the same system with large-scale applications — such as databases — that benefit from and have been rewritten for 64-bit addressability.

Additional factors that Aberdeen expects will play a role in the rapidity with which an ISV ports applications to Itanium are the size of the ISV and the number of platforms on which the ISV's applications already run. ISVs like SAS Institute, whose software already runs on a large number of platforms including 64-bit platforms, have well-defined methodologies and tools to simplify the porting of applications to new platforms. An ISV porting software to a new platform also incurs incremental costs — including not only the porting itself, but also testing, certification, and support. Nevertheless, Aberdeen expects that ISVs for whom porting software

to multiple platforms is a routine part of their software development process will find developing an Itanium port a straightforward process.

By contrast, software vendors whose applications currently run on a more limited set of platforms, especially those platforms that are 32-bit only, are likely to approach Itanium ports more conservatively. From the perspective of an IS buyer, Itanium versions of applications from such vendors are likely to be slower in coming to the market, becoming available only when driven by compelling business demands.

From a hardware perspective, Aberdeen expects that the first Itanium-based systems available to end-users will be quad-processor platforms based on the Intel 460GX chip set. These quad-processor systems will be based both on Intel's own Itanium-based motherboard and motherboards from leading system suppliers such as Compaq and Dell. Unisys, NEC, and Hitachi plan to offer larger 16- and 32-processor Itanium-based systems as well. Some of these platforms will leverage current designs such as the Unisys ES7000 Cellular Multiprocessing (CMP) system. Nevertheless, Aberdeen anticipates that the additional complexities and qualifications associated with large-scale system architectures will delay volume shipments of these high-processor-count systems until well into 2001.

### **How Itanium Will Perform on Key Application Types**

The combination of a 64-bit flat memory space, hardware resources such as a large register set, and a new floating-point architecture positions Itanium to deliver the greatest benefits to specific classes of applications. The delivery of these benefits generally assumes that applications are rewritten for 64-bit memory addressability if they are not already 64 bits and that the applications are recompiled for the Itanium architecture instruction set. As discussed in the section, "What IS Executives Need to Do to Migrate to Itanium" (page 18), IS managers have additional options that may require less effort when moving applications to Itanium-based platforms. ISVs may also take advantage of these lower effort options as part of a broader-based

Itanium deployment strategy in the future — especially for applications that can benefit from Itanium performance enhancements but do not require VLM. However, Aberdeen expects initial ISV applications offered for Itanium to focus on optimized application sets that leverage specific architectural features of Itanium with an emphasis on performance including, but not limited to, floating-point and large memory.

The Itanium architecture and its specific implementation in the form of the first Itanium processor have the opportunity to deliver competitive or leading metrics — even relative to 64-bit RISC processors — in key application and software infrastructure areas, particularly the following:

- Databases;

- Business Processing and e-Business;
- Business Intelligence;
- Technical and Scientific; and
- Security.

### *Databases*

Of all commercial application types, databases are those that typically drive the largest memory configurations. For example, IA-32 databases, such as Oracle, implemented paging extensions to allow the addressing of memory above the 32-bit (4 GB) logical address space but within the 36-bit (64 GB) physical addressability of the IA-32 processors. However, these paging extensions deliver significantly reduced performance compared to a 64-bit flat memory space in which all physical memory is addressed directly. Aberdeen notes that when RISC system vendors such as Sun, IBM, and Hewlett-Packard introduced 64-bit RISC processors, database vendors were among the first to port their applications to 64-bits to take advantage of the increased address range of the, then new, 64-bit processors.

Traditional databases such as Microsoft SQL Server and Oracle 8i Database Server avoid slow accesses to disk by buffering disk-resident data within main memory. As database sizes have grown rapidly, fueled in large part by high Internet-driven transaction rates, the need for larger memory buffers has grown correspondingly. The management of these large buffers has, in turn, generated a requirement for increased processing power to handle tasks such as background write daemons, read-ahead processing, and data partitioning. Microsoft SQL Server and Oracle 8i, among other databases, have been ported to Itanium and, when run on Itanium-based systems, will leverage extended memory addressability for larger buffering as well as the processing power to manage the buffers.

An emerging type of database takes a more direct approach to the use of the large quantities of physical memory enabled by 64-bit computer architectures. In-memory databases (IMDBs) are optimized for main-memory data management rather than disk-based management. TimesTen offers one such IMDB, which operates by offering direct pointers to records in memory. TimesTen reports that it implements database algorithms using only 10% of the instructions needed to do an equivalent amount of work in a disk-optimized database.

In addition to their performance and memory demands, databases are typically among the most mission-critical components of an enterprise's software suite. Such mission-critical database applications stand to benefit from advanced Itanium architecture reliability features that enable more reliable systems by detecting errors and minimizing the scope of their effects.

*Business Processing and e-Business*

The "Business Processing and e-Business" category covers a range of line of business (LOB) applications, ranging from traditional ERP like SAP R/3, to recently developed e-Commerce stacks from vendors such as CommerceOne and Broad-Vision. The common thread running through this broad range of applications addressing a wide variety of business needs is that the applications are large, complex, and are core components of an enterprise's business operations.

A summary of some of the applications planned for deployment on Itanium gives a sense of the breadth of this application category:

- SAP R/3 — traditional ERP;
- I2 RHYTHM — SCM, Customer Relationship Management (CRM), and service assets management for e-Business;
- SAP MySAP — an integrated e-Business platform evolved from SAP R/3;
- BroadVision — a line of Internet software products focused on personalizing e-Business;
- CommerceOne — e-Business products focused on business-to-business (B-to-B) services such as e-Procurement;
- Selectica Internet Selling System (ISS) — an enterprisewide application suite helping companies sell complex products over the Internet and through traditional sales channels; and
- Brokat Twister — an e-Services platform providing a technical platform and infrastructure for electronic commerce.

The large size and scope of many LOB applications underlies their ability to benefit from Itanium's large memory support. Of equal importance, however, are Itanium's instruction pipelining and large register set that allow more efficient processing of small or predictably sized integer data — common to many query operations — thereby increasing performance in many key LOB-application operations. Finally, the reliability enhancements made to the Itanium processor that handle error detection, correction, and containment can be used to help build the high-reliability system platforms required for core business software.

Users should also note that Itanium provides features applicable to many Internet applications including security (e.g., encryption algorithm support) and streaming media (e.g., floating-point performance improvements and support for applying a single instruction to multiple larger scale data streams). Floating-point and security performance are discussed in greater detail in the "Technical and Scientific" and "Security" application areas later in this section.

Aberdeen also observes that many LOB applications, especially in e-Business areas, were developed relatively recently and, in many cases, are already running on 64-bit RISC/Unix platforms. As a result, many of these applications lend themselves to

low-effort ports to a new platform such as one based on Itanium processors. However, Aberdeen expects that well-defined enterprise stacks like SAP R/3 will lead newer e-Business stacks into field deployment. This expected deployment timing is primarily a function of the complex software ecosystem required by typical large-scale e-Commerce and e-Business stacks to enable communication to take place with middleware, directory servers, security products, cache servers, and Web servers.

### *Business Intelligence*

The Itanium architecture's VLM support greatly benefits multi-terabyte data warehousing and data-mining installations. At the same time, Itanium architectural features, such as predication, that enhance ILP help ensure that even complex transactions with unpredictable control flows do not slow the bulk-load/process/bulk-store operations typical of large-scale and complex queries. The Itanium architecture's greatly increased floating-point performance relative to IA-32 is also significant for software performing complex numerical analyses of large data sets. Finally, the generation of explicitly parallel code for Itanium processors is of particular benefit to complex analytic software. The algorithms in such software are well understood by their developers who are also accustomed to tuning their code for specific processor architectures. By mapping code onto explicitly specified Itanium instruction bundles, developers of software with highly tuned algorithms can exert tight control over the execution of their code.

Vendors of business intelligence (BI) software tell Aberdeen that Itanium is highly significant to their applications and their customer base. These vendors see the Itanium architecture allowing more people access to high-performance hardware on which to perform increasingly complex analytics and time-series analysis on the huge volumes of data gathered from sources such as business-to-consumer (B-to-C) e-Commerce sites.

Products from SAS Institute, in addition to IBM Content Manager and Oracle Express, are among those in this category committed to deployment on Itanium-based systems.

### *Technical and Scientific*

The technical and scientific applications to which Itanium can bring benefits can be viewed as belonging to five classes:

1. Electronic Design Automation (EDA) — enables faster analysis of complex designs and headroom for complete data-set analysis;
2. Scientific — enables more complex analysis of fluid and chemical models;

3. Mechanical Design Automation (MDA) — manages large and complex assembly models and offers strong floating-point performance;
4. Digital Content Creation (DCC) — deals with larger and more complex models and scenes while enabling faster rendering at high resolution for 70-mm film and high-definition television (HDTV); and
5. Financial services — performs financial simulations in less time and delivers improved computational performance to handle industry changes, such as decimal trading.

The primary end-user purchase criterion for technical and scientific applications is usually performance. Intel estimates that an 800-MHz Itanium processor will deliver approximately 70% better performance on the Linpack-1000 benchmark — a standard benchmark for raw floating-point performance — than a 1-GHz Alpha 21264 processor, one of the highest-performing RISC processors. Based on this benchmark, applications ported to Itanium that make intensive use of floating-point would deliver performance on an Itanium-based platform as good as, or better than, that delivered by a comparable RISC platform.

Itanium architecture floating-point features, 64-bit addressability for large multimedia data-sets, and increased parallelism for the complex processes typical of technical workstations combine to deliver new performance and scalability benefits to a market already highly receptive to the price/performance delivered by Intel architectures. Itanium architecture floating-point registers will also provide better performance for many graphical applications.

Examples of technical and scientific applications include Fluent flow and heat transfer modeling software, Parametric Technology's CADD 5i computer-aided design (CAD)/computer-aided-manufacturing (CAM) suite, Alias|Wavefront's Studio 9 family of products for automotive design and styling, and risk-management software from the Riskmetrics Group.

### *Security*

Security is an increasingly important element of many electronic business transactions. However, the performance of encryption and decryption operations limits the scope of security-system deployment. Encrypting all network traffic from a client or a server requires significant processor resources. Web servers running the secure sockets layer (SSL) encryption protocol — used for receiving page requests and transmitting encrypted Web pages — face a 10-times to 20-times performance hit relative to Web servers that are not running SSL. Servers that handle multiple virtual private networking (VPN) sessions face a similar bottleneck because every TCP/IP packet sent and received by the server must be encrypted or decrypted.

Intel estimates that a production Itanium processor will perform 1024-bit RSA (Rivest-Shamir-Adleman) decryptions at approximately eight times the rate of a Sun

UltraSPARC III while also significantly exceeding the performance of an nFast security hardware accelerator. It is premature to suggest that Itanium-based systems will not benefit from hardware-based security acceleration. However, the speedy security performance expected from production Itanium processors is likely to benefit a broad spectrum of e-Business applications that make some use of security protocols, but are not dedicated to security operations.

In addition, 64-bit arithmetic also aids Itanium security-algorithm performance because these types of algorithms take blocks of data and perform operations on those blocks using a key. Because many of the blocks and keys are larger than 32-bits, 64-bit arithmetic operations allow public and private key indexes to be moved and manipulated in fewer steps. Other Itanium architecture features that help to improve security algorithm performance include loop-handling optimizations, a very large register set that stores intermediate results in complex algorithms, and the ability to execute significant portions of public key algorithms in parallel.

Planned security products for early deployment on Itanium-based servers include RSA CryptoC and Check Point Firewall.

### **Itanium Technology Adoption Life Cycle**

Significant insights into the adoption of the Itanium architecture and Itanium processors can be gained by mapping product and technology deliverables from Intel and its partners to the stages of the standard technology adoption life cycle. The groups within this life cycle are distinguished by their characteristic responses to a new technology:

- Innovators — the small group that first appreciates how the architecture of a product delivers competitive advantage to the marketplace;
- Early Adopters — the larger, but still non-mainstream, group that looks to new technologies to deliver fundamental breakthroughs in the way business is done;
- Early Majority — the largest part of the market for a technology product traditionally made up of pragmatic IS buyers looking to make incremental improvements to mission-critical systems while minimizing risk to the company's core business;
- Late Majority — the more conservative half of the mainstream market; and
- Laggards — the trailing group that resists new technologies and change.

From Aberdeen's perspective, the broad outlines of early-market technology adoption life cycle groupings align strongly with product and other deliverables in the rollout of Itanium. These same alignments point to challenges that will face both Intel and IS buyers as they move Itanium processor product sets from the laboratory to the data center. According to Aberdeen research, Table 1 lists the adoption of Itanium processor technology and products as viewed in this context.

**Table 1: Itanium Technology Adoption Life Cycle**

Buyer Category	Product Readiness	Buyer Challenges	Intel Challenges
Innovators	First production Itanium processor release  Targeted end-user pilots	Limited hardware and software infrastructure  Limited software tuning	Delivery of processors to enable end-user pilots  Translating technology wins into business advantage messages
Early Adopters	Production-qualified Itanium-based hardware and software  Expanded solution stack availability  Initial volume deployments	Developing, but still incomplete, hardware and software solution stack  Matching Itanium processor strengths to specific business problems	Enable targeted early adopters to achieve breakthrough business advantages  Continuing to enhance performance, reliability, and completeness
Early Majority	McKinley processor and supporting infrastructure availability  High volume deployments	Move from pilots and limited-volume deployments to large-scale adoption of Intel 64-bit  Managing 32-bit to 64-bit migration for maximum business benefit	"Crossing the chasm" to mainstream deployment  Delivery of McKinley and complete hardware and software solution stack

Source: Aberdeen Group, November 2000

**Early Adopter Profiles**

Early adopters of the Itanium processor face the key challenge of properly matching the strengths of the processor and architectural infrastructure to specific business problems. In Aberdeen's view, the Itanium processor can deliver significant enhancements to critical business processes with unmet computational needs that align with Itanium processor strengths and well-defined software stacks. The following profiles of candidates for early adoption of the Itanium architecture are not intended to comprehensively represent all IS buyers that could potentially benefit from the Itanium processor. Nor should the following profiles be taken to suggest that other buyers should not initiate pilots or evaluations with an eye on later production deployment. Rather, these profiles represent the characteristics of those specific IS-buyer characteristics that Aberdeen expects to be among those that can derive the greatest benefit from the earliest Itanium processor production deployment.



*Pushing Database Performance to the Next Level*

This first group of early adopters has been chartered with revamping a legacy database system that is increasingly incapable of handling the peak transaction loads generated by e-Commerce applications. Planned integration of additional B-to-C applications that bypass traditional transaction intermediaries is expected to dramatically increase volumes further over the next one year to two years. The bottom line: Incremental increases in transaction capacity will be insufficient to handle projected capacity needs.

IMDBs are the most direct approach to addressing critical database performance bottlenecks. Holding an entire database in memory can either be implemented in the form of a newer, specialized database like TimesTen — designed specifically for that purpose — or in the form of large buffers for a more traditional database such as Oracle 8i.

In many application architectures, databases already exist quasi-independently of other system components. For example, in many Internet environments, databases are hosted on large Unix servers while other elements of the compute environment run on commodity servers running Microsoft Windows or Linux. Furthermore, databases typically have well-defined interfaces into other applications whether directly or through middleware components. These well-defined interfaces allow databases, in many cases, to be integrated into an application environment with minimal disruption to other components.

Early adopters considering Itanium-based systems for IMDBs will also consider large servers from RISC/Unix vendors such as Sun. Those who settle on Itanium-based systems will be attracted to the long life they anticipate from Intel's new Itanium architecture and the potential for commodity hardware pricing. Alternatively, early buyers of Itanium-based systems for large-scale databases may have an eye on current or future deployment of Windows 2000 at the back-end of their enterprises. As Aberdeen noted in the research report, *Is Microsoft's Windows NT/Windows 2000 Enterprise-Ready?* (January 2000), many IS buyers are gaining increasing confidence in Microsoft operating-system environments as a foundation for mission-critical environments, including databases.

*Expanding the Scope of Business Intelligence Deployments*

Rapidly increasing quantities of data affect not only the ability of back-end database systems to process and store data, but also the ability of IS executives to wade through the data for the purpose of developing quantitative and qualitative business insights. This second group of early adopters wants to distribute deep analysis of core business data across a broader range of functional units in the enterprise.

BI software like that from the SAS Institute already runs across a range of 64-bit and 32-bit hardware and OS platforms. However, the commodity hardware pricing trend, which Intel processors have helped to drive in 32-bit architectures, is likely to repeat with Intel Itanium processors. This lower pricing will aid in the more widespread distribution of Itanium-based hardware throughout an enterprise to those applications, such as exchange, that can immediately benefit from both the performance and the large memory features provided by Itanium. By distributing the ability to gain business insights through quantitative analysis to the functional units with the need for the insights, early adopters of BI applications on Itanium processors hope to enable groups within their organizations to make more informed business decisions more quickly.

BI applications have differing degrees of dependence on other components of a software stack. At a minimum, interfaces to a back-end database are required in some form, typically through middleware that also performs various data-cleansing and transformation functions.

Early adopters considering Itanium-based systems for BI applications will also consider 64-bit RISC/Unix servers and desktops. Pricing is likely to be an Itanium advantage — at least over time. An additional Itanium benefit could include performance, especially as algorithms are tuned to the explicitly parallel Itanium architecture over time.

### *Unifying Business and Technical Computing on Intel Architectures*

Technical and scientific computing deployments on Intel architecture desktops and servers have been limited in the past by IA-32 floating-point performance relative to RISC architectures. As a result of this past floating-point disparity between IA-32 and RISC, many enterprise computing environments use Intel architecture systems for commercial applications while using RISC/Unix systems for their technical applications. This group of early adopters has a vision of reducing costs and management effort by standardizing on an Intel architecture platform across both commercial and technical applications.

This third group of early adopters may also be intrigued either by the possibility of migrating to a pure Windows 2000 environment across both business and technical application types in their enterprise or by the possibility of adopting Linux to deliver high levels of compute performance on desktops or small servers at commodity pricing levels. Aberdeen expects that Linux is likely to be deployed by a broad range of Itanium early adopters in the scientific and technical application spaces. This early Linux deployment on Itanium-based systems will be driven by the innovator and early adopter characteristics shared by many users of both Linux and technical/scientific applications. By contrast, Windows 2000 deployments are likelier to mirror mainstream buying characteristics and timing.

Among the different categories of technical computing, the porting and tuning of code for Itanium and planning for Itanium-based application pilots is furthest along in DCC, MDA, and pure scientific applications. Therefore, it is these specific application segments that are best suited for initial deployments of Itanium-based systems with other application types throughout 2001.

Performance that exceeds available RISC architectures would provide a compelling impetus to early adopters considering Itanium processors in the technical and scientific computing space. However, performance parity with leading RISC/Unix processors — if mated with the availability of the necessary software — would lead many IS buyers to consider standardizing on Intel architectures as a means of simplifying their computing infrastructures.

### **What IS Executives Need to Do to Migrate to Itanium**

While most initial Itanium-based system deployments will be for new applications often based on third-party software products, Aberdeen recognizes that different IS requirements will drive a variety of application migrations. To help identify the correct migration strategy and time frame for in-house computer code, Aberdeen recommends that all IS executives planning to migrate to Itanium systems take the following steps to prepare:

1. Evaluate which in-house applications will benefit from the Itanium architecture and Itanium processors and the OSs they run on — are both ready for Itanium? If not, when will they be ready?
2. If the customer has a choice among suppliers for Itanium-based systems, evaluate which system will require less migration effort; which will allow the user to take more immediate advantage of Itanium processor performance and reliability features; and which system supplier provides greater support for the migration effort. The key is to begin asking system providers today about their strategic Itanium plans and to identify those organizations with the greatest depth of knowledge and the most strategic commitment.
3. For key performance-critical applications, begin to test the effects of the Itanium processor using compilers — whether provided by the supplier or a third-party. Users forging partnerships with system suppliers may want to arrange to take advantage of these suppliers' own testing environments.
4. Create a plan for migration, including specification of testing systems and a time line for moving migrated applications into production. IS managers should consider both third-party tools and services to aid them in determining software dependencies.

5. Move selected applications to an Itanium-based platform with key performance-critical or resource-intensive applications rewritten to take advantage of Itanium and other applications simply recompiled. The running of IA-32 binaries on Itanium-based platforms should be minimized and phased out over time.

Once enterprises have identified targets for migration, IS executives should also, for planning purposes, evaluate when it makes sense to migrate to Itanium-based systems, in whole or in part, and periodically make re-evaluations based both on internal business needs and on external product delivery schedules. For most IS executives, migration will involve either a recompile or a simple port for most internal applications and an upgrade from the ISV for packaged applications. The good news for users of packaged applications is that many ISVs are being proactive in delivering versions of their products that support Itanium; however, IS managers should, of course, check with all ISVs that compose their application environment.

The effort to migrate or port internally developed applications will vary considerably depending on coding practices, complexity, and the degree to which software modules and libraries interact. Before a porting project can begin, it is necessary to locate all affected source code and header files, and third-party libraries along with linkages to and dependencies on other applications. The effort associated with this stage of the porting process may be small or nonexistent in well-organized environments — especially among actively maintained applications with up-to-date build environments — but can be significant for some legacy code.

Aberdeen notes that vendors such as Hewlett-Packard and Microsoft include tools in their development kits to augment base compiler functions in the identification and remediation of problem code. Third-party tools like the MigraTEC Migration-SUITE may provide additional assistance to IS managers responsible for porting code in-house by not only providing significant automation of code porting, but also helping these managers to better understand their current code environments. Other IS managers may choose to use third-party services for porting.

Beyond common issues that affect code porting across most environments, certain specific migration details and methodologies will vary by the type of system presently used:

- Windows NT/2000;
- Intel 32-bit Unix and Linux; or
- RISC 64-bit Unix.

#### *Windows NT/2000*

Microsoft describes Windows 2000 64-bit edition (Win64) as a “logical widening” of 32-bit Windows 2000 that adds new data types but largely preserves the 32-bit Windows application programming interface (API). These new data types are a combi-

nation of new explicit-precision 64-bit pointers and data types, and architecturally neutral data types that will be sized appropriately depending on the platform.

Win64 provides for a single source base for 32-bit and 64-bit software products running under Windows 2000 by allowing Windows APIs to float their precision to either a 32-bit or a 64-bit platform. The new architecturally neutral data types are now also part of the Microsoft Windows 2000 32-bit development environment to provide this compatibility.

Microsoft indicates that many IA-32 applications written in a portable manner should require only a recompile to run on Itanium processors using the native Itanium architecture instruction set. Thus, migrating a Windows-based application should require a recompiled version of the application and a system upgrade to Win64. However, Aberdeen recommends that IS managers with mission-critical Windows applications planned for Itanium-based environments consider making the code changes required for full 64-bit support. In this way, the applications can take advantage of VLM now or in the future, and incremental port, debug, and test cycles can be avoided. By bringing applications to full 64-bit readiness, IS managers can also avoid certain potential software incompatibilities such as those that can occur in dynamic link library (DLL) interactions between 32-bit and 64-bit application components.

#### *Intel 32-Bit Unix and Linux*

The IBM and SCO AIX 5L effort, formerly Monterey/64, is available in beta form as are four separate distributions of the developer release from the IA-64 Linux project. In the case of AIX 5L, two primary programming modes are available: 32-bit on the Itanium architecture instruction set (ILP32) and full 64-bit (LP64). IA-64 Linux is available with only the full 64-bit programming model. The 32-bit on 64-bit model may require little more than a recompile of the application to run on Itanium-based systems; the full 64-bit model will also require some amount of re-writing of legacy applications — especially those that make casts between unlike data types whose sizes change in the move to 64-bits. The issues are similar to those in Windows environments although some of the details are different because of different 64-bit programming models chosen for Windows and Linux/Unix. For example, long data types are 64-bit on Unix and 32-bit on Win64. For many applications, a relatively straightforward recompile of 32-bit code to the Itanium architecture instruction set will likely be sufficient. Aberdeen recommends that performance-critical applications or those with significant memory needs — even if they do not require 64-bit addressability today — be rewritten to full 64-bit.

#### *RISC 64-bit Unix*

Suppliers such as IBM (AIX) and HP (HP-UX) aim to offer Unix operating systems on Itanium-based systems and to provide extensive services for users of alternative

64-bit architectures looking to migrate to Itanium-based systems. These suppliers also provide specific tools to migrate existing applications on 32-bit systems to Itanium-based systems. For IS executives, migration may involve moving to the supplier's version of Unix as well as from 32-bit to 64-bit systems. In many of these cases, applications will not require extensive rewriting, assuming that they were written in a portable manner. Code that has been written to more tightly integrate with a given vendor's flavor of Unix will require more extensive recoding.

#### *How the Itanium Architecture Addresses 32-bit/64-bit Binary Compatibility*

Itanium's architects have designed Itanium to provide backward binary compatibility with IA-32 applications that run on Intel's X86 platform. A portion of Itanium processor real estate is dedicated to processing 32-bit code, and the Itanium architecture system environment allows either IA-32 or Itanium architecture instructions to be executed at any given point in time. Software support for IA-32 binary execution on an Itanium-based platform rests with the operating system; however, all announced operating systems for Itanium plan to support IA-32 binary execution. Aberdeen cautions that Itanium binaries and IA-32 binaries will typically execute in separate process spaces, and therefore data and file sharing will be limited.

The IA-32 compatibility provided by the Itanium architecture should be considered solely for compatibility, not performance. IS executives looking for maximum performance for an application will require applications recompiled for Itanium. For maximum IA-32 performance, IS executives should strongly consider Intel IA-32, rather than Itanium processors. As a rule of thumb, Aberdeen suggests IS managers should aim for at least 75% of the total code and all of the performance-critical code on an Itanium processor to have been recompiled — if not rewritten — for Itanium.

#### **The Competitive Landscape and Challenges for Itanium**

In spite of the availability of some early benchmarks, it is too early to clearly conclude how Itanium processors will perform quantitatively against other processor architectures on the market. Intel claims that the Itanium architecture will be more than competitive against existing enterprise RISC architectures. However, it is clear that there will be a significant number of alternative architectures available. As a result, it will take time for the Itanium architecture to gain momentum.

Among other competitors, IBM will continue to push the performance envelope with its POWER microprocessor architecture, making headlines with its Silicon-On-Insulator and copper CPU-technology innovations. Sun continues to gain market share in the commercial application Unix-server market, especially in Internet servers with successive generations of its UltraSPARC — most recently the UltraSPARC III. Hewlett-Packard's PA-RISC processor roadmap spans three additional generations from the PA-8600, on which the HP 9000 Superdome server is initially

based, through the PA-8800. At least one additional generation is planned beyond the PA-8800. Compaq continues to design new generations of its Alpha micro-processor and expects to deliver copper-based EV68 parts that may reach 1250 GHz by mid-2001.

Despite the competition, Itanium has seen support from all of the aforementioned system vendors, with the exception of Sun. In fact, several vendors have committed to discontinue their RISC processor lines — over time — in favor of Itanium processors. Aberdeen notes that IBM and HP have announced plans to run their Unix operating environments on Itanium-based systems — a factor that will contribute to the adoption of Itanium in the high-end commercial server market segment.

Another potential competitor for Itanium down the road is Advanced Micro Devices' (AMD's) forthcoming line of 64-bit processors (code-named Sledgehammer) for which AMD released technical information for developers in August 2000. The design approach taken by AMD with these processors is to enhance the current x86 processor architecture in an evolutionary manner — similar to the approach taken when Intel processors moved from 16-bits to 32-bits — rather than pursue a radical redesign as Intel has done with the Itanium architecture. AMD's approach has the probable advantage of superior 32-bit x86 binary performance relative to 32-bit binaries running on Itanium-based systems. It remains to be seen, however, how 64-bit performance will compare with the AMD and Itanium processor families over time and also the degree of ISV commitment that AMD is able to obtain for its 64-bit architecture. It is currently anticipated that AMD 64-bit processors will be available in a similar time frame to Intel's McKinley — late 2001.

### **Aberdeen Conclusions**

The need for more computing power spirals upward as data-intensive decision support, Internet-based multimedia, and transaction-oriented applications gobble up as much computing power as an enterprise can put forward. Intel's Itanium architecture is expected to help IS executives keep pace with user demands for more power and better response times based on both specific Itanium processor implementations and the underlying design philosophy of EPIC.

But Aberdeen suggests that hot technology is not enough to drive the success of Itanium. Itanium industry support will be the key factor in its success. OEM and OSV support has, so far, been significant. Intel and its partners such as Hewlett-Packard, IBM, Unisys, and many others have all been working aggressively with their ISVs to ensure that complete solution suites are available for Itanium-based system production. And, by doing so, the systems vendors and Intel may be taking much of the transitional effort of moving from 32-bit to 64-bit computing off the shoulders of the IS department.

Widespread deployment of the first Itanium processor should be viewed as a strategic move for specific early-adopter IT departments with profiles aligning with

Itanium capabilities and software solution stack availability. Aberdeen has identified a specific set of profiles that are expected to benefit disproportionately from early Itanium-based system deployment:

- Pushing database performance to the next level;
- Expanding the scope of business intelligence deployments; and
- Unifying business and technical computing on Intel architectures.

Aberdeen also suggests that a much broader swath of IS buyers should evaluate the use of Itanium-based systems, deploying pilots where appropriate, across a wide range of new application environments, including e-Business. These new applications are less likely to have complex interactions with legacy systems and will give IT the opportunity to evaluate Itanium's potential for broader deployments in their enterprises.

If the Itanium architecture captures the industry's leading ISV applications, the high end of the computing market could experience the same "commoditization" effect at the processor level (and then systems level) that the industry saw when Intel moved into the PC marketplace and the low-end/midrange server marketplace. Volume manufacturing tends to drive costs, resulting in an overall lowering of computer hardware costs as well as increased competition on the performance side of the equation — and the lower costs/higher performance dynamic has been beneficial for the IS buyer.

Challenges for Itanium remain. The greatest challenge is eliminating or minimizing further slips in the rollout of the first Itanium processor — fulfilling a schedule that requires consensus among the application suppliers, OS suppliers, and system suppliers to build a strikingly different enterprise platform. The next challenge is making sure that EPIC compilers are optimized to handle key applications and work closely with the microprocessor to provide the compelling performance that IS executives demand in enterprise environments. The third and final hurdle will be to continually grow interest and development among packaged-application ISVs that have not yet committed to Itanium.

The initial Itanium processor is a necessary first step in a new computing architecture with the potential to cause significant shifts in enterprise system deployment over the next decade. Just because it is a first step, however, does not mean that it should be relegated to evaluations or pilot deployments. IS buyers with relatively static compute environments or without a near-term application match to Itanium architecture capabilities, such as large memory, may indeed want to move cautiously onto Itanium-based platforms. Early adopters with the right profiles, however, could find in Itanium the lever that helps deliver breakthrough business advantage as part of a new type or style of application deployment in their enterprises.



## Appendix

### **Itanium Architecture Overview**

EPIC relies heavily on processor-savvy compilers to optimize performance. The need for tight integration with compilers means that Intel is committing significant resources to line up key industry players (system suppliers, ISVs, and services providers) so that the initial Itanium processor and future products from the Itanium processor family can provide comprehensive performance solutions “out of the starting gate.” The Itanium architecture also aims to achieve a new level of parallel processing at the instruction level. New parallel-processing features include hardware support for multiple parallel instruction streams, lower memory latencies than traditional architectures, explicit specification of parallel behavior at compile time, an increase in the number of parallel streams supported, and a third level of caching.

Aberdeen concludes that Itanium processors will change the way IS executives look at Intel architecture platforms. For example, most IS executives have used frequency (i.e., megahertz) on the Intel architecture as one method of determining performance. With the introduction of the Itanium architecture, as with other new computer architectures, it becomes increasingly important for IS executives to measure total system performance based on a number of factors that include other elements such as server I/O, memory, the instruction set itself, and key application performance.

### **Itanium Architecture Key Components: Breaking New Ground**

The philosophy underlying EPIC is to enable processors to handle more instructions per clock cycle (IPC) — more application code — and to feed these instructions to multiple on-chip functional units for execution in every clock cycle. Intel's Itanium architecture designers are building significant performance headroom into the architecture, which allows Intel to cram in more IPC over the next few generations of the Itanium processor family.

In traditional architectures, a key limit to performance and scalability has been the ability of the architecture to keep the multiple execution units in the processor fed with instructions and data. In most of today's microprocessor architectures, the processor has the task of dispatching instructions to execution units on the fly — a task that consumes valuable real estate and processing power on the processor. This dispatching is also difficult for the processor to do efficiently in real time — with no knowledge of the underlying code. By building more and more execution units, current RISC and CISC (complex instruction set computer) architectures have been able to continually boost performance. But it has become increasingly

difficult for these architectures to fully utilize the execution units — leaving an increasingly large potential performance boost on the table.

*The Itanium Architecture's Crown Jewel: Application Code Processing Efficiency*

With the Itanium architecture, the compiler takes the bulk of the responsibility for creating parallel application code. The compiler packages parallel application code specifically for the Itanium architecture before it is sent to the processor. In effect, the compiler gives the processor the blueprint for a parallel instruction stream embedded in the instruction stream itself, and the processor carries out that blueprint, adding its own optimizations. This process is particularly effective where the big opportunities for performance improvements are in optimizing large chunks of code for parallel execution, because processors cannot “look ahead” very far in the instruction stream to detect these code chunks.

It could take some time for the new compiler-driven optimization to translate into major jumps in application-level performance. History has shown that compilers' programmers take several years to learn how to optimize most effectively. However, compiler vendors — including hardware suppliers, Microsoft, and Edinburgh Portable Computers (EPC) — will be two years or more into the learning process when Itanium arrives. Key ISVs have told Aberdeen that the overall level of maturity of compilers and other tools for Itanium is quite good. Nevertheless, IS executives expect continuing performance improvements over time from Itanium compilers and applications — especially in light of the greatly increased portion of the optimization task that compilers take on in the Itanium environment.

Intel has introduced the following key technologies to boost the number of instructions that the processor can handle and a number of other key technologies that build on EPIC. These technologies include:

- Predication;
- Speculation;
- Rotating registers and other processing efficiencies; and
- Hardware enhancements (larger caches, enhanced scalability, and larger integer and floating point units).

*Predication*

A key factor in the ability of a processor to maximize instruction-level parallelism is following the correct branch at decision points in the code. Traditional architectures cannot easily see far ahead in the instruction stream, and therefore do not know which instruction stream after the decision point is the correct one. To avoid predicting incorrectly which branch/instruction stream to execute, today's processors may attempt to guess more intelligently (“branch prediction”), run both instruction streams in different pipelines (“speculative execution”), or exe-

cute one or both of the branches when the processor is underutilized (“out-of-order execution”). Much of the time, these tactics do not fully leverage processor parallelism: Intel estimates that these architectures’ “mispredicted branches” (executing the wrong instruction stream) can cost up to 40% in processor performance.

Predication takes a new approach to the problem — i.e., “looking ahead” in the compiler before the processor receives the code. Predication removes branches wherever possible. By removing branches in this manner, larger contiguous chunks of code are fed into the highly pipelined Itanium processors. Pipelining can, on the one hand, enable high levels of performance by breaking instruction execution into multiple autonomous operations, which can be executed rapidly and consecutively. On the other hand, as pipelines increase in length, the performance penalties associated with flushing a pipeline — effectively forcing the restart of a whole train of instructions — increase correspondingly. By minimizing the degree to which code is broken into small segments by branches, with consequences such as pipeline flushes, predication improves overall processor utilization. Applications involving complex, frequently branching code — legacy, technical/scientific, and computation-bound solutions like economic order-quantity computation in ERP, stock-market analysis, and data analysis for data mining — will particularly benefit from predication.

### *Speculation*

Instructions and data that arrive at the processor late (from main memory or cache) slow performance. Yet, although the processor may be waiting for data from memory in some portions of an application, the memory subsystem may be sitting idle with nothing to do in other portions.

Itanium architecture speculation technology loads instructions from main memory or cache before the processor is ready for them — even if it turns out that the processor will not need to execute the instruction. In effect, speculation caches instructions and data on the processor itself. The difficult aspect of this technique used by the Itanium architecture is that memory load instructions can sometimes cause software exceptions or errors that must be handled properly by the executing program at the correct location in the code. The Itanium architecture solves this problem by introducing instructions and registers that separate load behavior from error-handling behavior. These devices allow load instructions to be “hoisted” above intervening branches.

Speculation is particularly effective with applications that access memory frequently, because these applications need to quickly feed instructions and data to the processor. For example, large-scale Internet or database applications must switch rapidly between large numbers of end-users; speculation can “pre-send” the

instructions and data for each new end-user to the processor so that the system can react more rapidly.

The Itanium architecture also can combine predication and speculation — eliminating many branches through predication — then speed execution of the resulting code through speculation.

#### *Rotating Registers and Other Processing Efficiencies*

The Itanium architecture supports 128 64-bit-wide registers that system programmers can access. These registers allow the processor to process more data at one time. Itanium architecture designers have added a number of monitoring and management features that will allow programmers to use these registers more effectively. For example, the Itanium architecture provides “rotating registers” that act like an assembly line, allowing the processor to apply an operation to each register in turn. This approach means that when the same set of operations is to be applied to a stream of data, the processor can “overlap” the operations for a high degree of parallelism. Key application functions such as large database sorts and technical/scientific floating-point calculations (a necessity for many of today’s workstation applications) will benefit from this feature.

The Itanium processor also includes a number of performance-monitoring registers. For example, certain key registers keep a count of instructions per minute and cache misses. It is likely that, in the future, systems-management software will be able to monitor these registers in real time, avoiding systems-management overhead and improving a user’s ability to “tune” processor performance to the workload.

#### *Other Hardware Enhancements*

For technical/scientific computing environments, Itanium provides higher precision floating-point and integer-point units that can more effectively handle the large/precise-number computations typical of these environments. The Itanium processor provides 128 floating-point registers in addition to 128 general registers, resulting in more write and read ports. For applications such as technical-workstation solutions and Internet-focused Digital Content Creation applications, the graphics quality, degree of rendering accuracy, and speed of processing will be enhanced. For the larger-data-size applications, typical of technical and scientific markets, Itanium will support two extended-precision-floating-point multiply and accumulate instructions, which are capable of single, double, and extended precision-floating-point calculations.

The Itanium processor also includes a multilevel bus and cache structure. Level 1, with separate address and data caches, and Level 2 caches are incorporated on the

processor die, and a large (4-MB) Level 3 cache is connected to the processor through a dedicated 128-bit bus running at full processor clock speed.

### **Itanium Architecture Intelligent Hardware Reliability Features**

The processor, firmware, and the Itanium-hosted OS work together to contain and fix errors that occur in application environments. The Itanium processor signals the error to the firmware, which can then either correct the error or supply information to the OS for software correction. This feature provides additional opportunities to correct errors, instead of crashing, resulting in greater system availability.

Another reliability enhancement, called data poisoning, is the flagging of double-bit ECC errors in the processor caches. Data poisoning allows the system to continue without a reboot; only the offending process needs to be terminated via machine check abort (MCA) when the poisoned data is consumed (i.e., when the poisoned data is accessed from the caches). A special processor state allows unaffected processes to continue instead of bringing down the entire system. Other features include error logging — giving IS administrators the ability to trace application, I/O device, OS, and transaction errors — so that IS administrators can determine where faults are occurring in the system.

And, while Itanium architecture designers have focused on processor-centric enhancements, other Intel and OEM hardware system designers have started to explore ways to eliminate reliability issues found in today's server I/O architectures. The InfiniBand initiative is working to introduce switched-fabric I/O-system backbones to augment and eventually replace server-based PCI bus architectures over time. System suppliers will also offer a number of other reliability features that have become common for enterprise servers such as redundant fans, PCI Hot Plug components, redundant power supplies, and other reliability features.

---

*Aberdeen Group, Inc.  
One Boston Place  
Boston, Massachusetts  
02108  
USA*

*Telephone: 617 723 7890  
Fax: 617 723 7897  
www.aberdeen.com*

*© 2000 Aberdeen Group, Inc.  
All rights reserved  
November 2000*

Aberdeen Group is a computer and communications research and consulting organization closely monitoring enterprise-user needs, technological changes and market developments.

Based on a comprehensive analytical framework, Aberdeen provides fresh insights into the future of computing and networking and the implications for users and the industry.

Aberdeen Group performs specific projects for a select group of domestic and international clients requiring strategic and tactical advice and hard answers on how to manage computer and communications technology.